

IMPROVING DIGITAL HALFTONES BY EXPLOITING VISUAL SYSTEM PROPERTIES

Jeffrey B. Mulligan

NDB

PSB-NASA

7N-53-CR

026074

Reprinted from CONFERENCE RECORD OF THE TWENTY-
SEVENTH CONFERENCE ON SIGNALS, SYSTEMS &
COMPUTERS, Pacific Grove, California, November 1-3, 1993

**Improving Digital Halftones by Exploiting
Visual System Properties**

By

Jeffrey B. Mulligan

Improving Digital Halftones by Exploiting Visual System Properties

Jeffrey B. Mulligan

NASA Ames Research Center, MS 262-2, Moffett Field CA 94035-1000

1. Abstract

The visibility of quantization noise in digital halftones can be predicted from psychological data on spatial and temporal contrast sensitivity. Simple models of the visual system can be incorporated into halftoning algorithms to minimize the visibility of the resulting artifacts. Filter-based algorithms may be customized to match the error filter to human contrast sensitivity under known viewing conditions. The relative insensitivity of the visual system to high frequency chromatic modulation allows visible luminance noise to be reduced at the expense of additional (but invisible!) chromatic noise. The techniques are easily extended to three dimensions for displays which can be modulated in time such as CRT's and flat panel displays.

2. Introduction

The problem of representing continuous tone images on a binary display device is known as "halftoning" or "dithering" (a good survey can be found in Ulichney [1]). Halftoning works because the human visual system integrates information over spatial regions, so that a spatial pattern of light and dark can evoke a sensation approximating that of a uniform gray area even when the individual display elements can be resolved. When the observer is far away, or the display device has extremely high resolution, the dithered version may be indistinguishable from the original. In many cases, however, the dithered image *will* be distinguishable from the original. The *quantization noise* is the difference image formed by subtracting the original image

from the dithered image. The goal of visually-optimized halftoning is to make the quantization noise as invisible as possible while maximizing the visual fidelity of the encoded image.

In this paper we will consider how existing dithering algorithms can be improved and/or tuned to particular viewing situations by the application of knowledge of parameters of the human visual system. Three important variables which are important to the visual system are space, time, and color. Knowledge of spatial sensitivity to visual patterns allows us to tune algorithms for optimal performance at a particular viewing distance. On dynamically controllable devices such as CRT's and flat panel displays, patterning in time can be used to increase gray level resolution and improve perceptual segregation of the signal image and the halftoning noise. When color images are processed, the different parameters describing spatial and temporal sensitivity to chromatic and achromatic image components can be exploited to improve the rendition of the achromatic component by shifting noise to chromatic bands to which observers are less sensitive.

3. Visual System Parameters

3.1. Spatial Sensitivity

The human visual system is not equally sensitive to spatial patterns of different sizes. The *contrast sensitivity function* (CSF) describes the visual response to spatial patterns as a function of spatial frequency [2,3]. The units used to describe spatial frequency depend on which journal one is reading. In the vision literature, spatial

frequency is usually given in cycles per degree of visual angle, while in image processing it is usually specified in cycles per image, or cycles per pixel. In order to equate these two scales, it is necessary to first define the viewing conditions. This dependence upon viewing distance makes it impossible to create a single visually optimal dithering algorithm; in some applications, the best choice may be a compromise set of parameters which produce optimal rendition at no viewing distance, but provide acceptable rendition over a broad range.

Dithering algorithms vary in the ease with which they may be spatially tuned. Ordered dither is fast and parallelizable, but is not particularly easy to tune. Mitsa [4] has proposed a method for generating "blue noise" threshold arrays which produce dither textures similar to those produced by the popular Floyd-Steinberg error diffusion algorithm [5]. This procedure requires storing a large threshold array which is not easily generated, however. Unlike the other algorithms to be described, the ordered dither algorithm processes pixels at different spatial locations independently; pixels are quantized based solely on their own values, without regard to the values of neighboring pixels.

The error diffusion algorithm [5] is often preferred to ordered dither because it does a better job of rendering small details, and does not generate quantization noise with a periodic structure. The algorithm quantizes the pixels sequentially, which is required because, unlike the ordered-dither algorithm, pixels at different locations are not quantized independently. As each pixel is processed, the quantization error is spread or "diffused" to neighboring pixels which have yet to be quantized. Successive pixels are quantized not solely on the basis of the corresponding values from the original input image, but rather from the values modified by the error from the quantized neighbors. In this way, neighboring pixels are quantized so as to produce cancelling errors.

The spatial quality of the resulting quantization noise is determined by the weights which are used to spread the error. These weights can be

thought of as the kernel of a low-pass filter which specifies the signal band of interest. In actual implementations the weights are often chosen to be rational fractions involving powers of two, allowing the computations to be done using fast integer arithmetic. If we are more concerned with visual quality than with speed of computation, we can tune the noise by adjusting the filter weights to produce a filter which mimics contrast sensitivity under particular viewing conditions. The biggest obstacle to doing this properly is that the sequential nature of the algorithm forces the filter to be causal, i.e. the errors can only be spread to the side of the pixel which has yet to be quantized.

This problem can be eliminated, but only at a large computational cost. A number of iterative algorithms have been proposed which seek the halftone image which minimizes the error filtered by an arbitrary filter [6-11]. When non-causal filters are used, changing the state of a given output pixel may necessitate changes to previously quantized neighbors; thus these algorithms typically require many passes over the image before a stable configuration is obtained. The complete freedom with which the filter may be designed, however, allows the algorithm to be precisely tuned to a particular application.

3.2. Temporal Sensitivity

The use of dynamically controllable displays such as CRT's and flat panel displays has created a new opportunity for the improvement of halftones through the use of *temporal dithering*. Temporal dithering refers to the rendition of a desired gray level with a spatial distribution of *flickering* pixels. The utility of this approach is rooted in the fact that the visual system's spatio-temporal sensitivity cannot be described as the separable product of a spatial CSF and a temporal CSF, but shows dramatically reduced sensitivity to high spatial frequencies when presented at high flicker rates, while sensitivity to low spatial frequencies is unaffected by flicker [2,3,12-16]. For low to medium spatial frequencies, temporal sensitivity is bandpass with peak sensitivity between 5-10 Hz, while for patterns above 4 cycles per degree temporal sensitivity becomes low-pass.

This feature of the visual system is exploited in the use of interlaced video signals: each video pixel flickers at 30 Hz, which is clearly visible when pixels are lit individually. When a large area is lit, however, it is a high spatial frequency pattern of light and dark lines which is modulated at 30 Hz, which *cannot* be seen, even though the visual system has sufficient spatial resolution to see this pattern when presented at lower temporal frequencies.

Mulligan [17] has proposed a number of ways in which existing spatial dithering algorithms may be generalized to three dimensions. Ordered dither is easily generalized by the construction of three dimensional threshold arrays. The principles by which these arrays should be constructed to produce optimal visual quality are not well understood, but it is straightforward enough to compare the performance of two arrays on a particular image if one has a computational model of the visibility of artifacts.

Methods involving error filtering are easily extended to three dimensions by simply filtering the error in the time domain as well as the space domain. In the case of error diffusion, this corresponds to diffusing a fraction of the error into the next temporal frame. Purely temporal error diffusion can be used to turn *any* spatial dithering algorithm into a hybrid space-time algorithm by applying the spatial algorithm to the first frame, computing the quantization error image, and then subtracting this error image from the next frame before spatial processing.

These methods produce image sequences which are superior to static halftones in a number of ways: first, better gray level resolution may be obtained, because more quantized pixels are averaged together by the visual system; secondly, and perhaps more importantly, the signal and noise are concentrated in different temporal bands, which greatly enhances the perceptual segregation of the two. The percept is sometimes evoked of a picture viewed behind a dirty window (with the dirt moving around dynamically). The idea is that it is better to have a lot of dirt on a window which you can choose to ignore, than a little bit of dirt stuck on the picture itself, which cannot be

segregated. Modeling this perceptual segregation is an open problem in vision science, which when solved will provide guidance for the next generation of halftoning algorithms.

3.3. Color

The response of the human visual system to color is markedly different than its response to achromatic or luminance information. For the purposes of dithering, the important facts are that the chromatic subsystem is low-pass both in space and time. For example, if a pattern of colored stripes is progressively minified, at some point the colors of the individual stripes will blend and the pattern will appear to have variations only of intensity, and the pattern will be completely invisible if the colors are made *equiluminant* (equal visual intensities). This makes it sensible to concentrate on minimizing high frequency luminance artifacts in halftones, while accepting large chromatic errors which will not be visible.

The differential *spatial* properties of the chromatic and achromatic systems have been exploited in several halftoning algorithms. Mulligan [18] proposed a color ordered dither algorithm in which negative correlations are introduced into the threshold matrices used to create the component subimages. Compared to using identical matrices for the component images, this has the effect of reducing luminance noise while increasing chromatic noise. Because of the different sensitivities to these features, this can often result in an improved image. Mulligan and Ahumada [11], and Balasubramanian, Carrara and Allebach [10] used different spatial error filters for achromatic and chromatic quantization errors in iterative algorithms designed to find the visually optimal halftone.

Additional benefits may be reaped from the *temporal* low-pass character of the chromatic system. When two colored lights are exchanged or flickered, the color will appear to alternate at low flicker rates, but when the frequency is raised to 15-20 Hz, *color flicker fusion* occurs, where a single steady color is seen and the flicker is seen as a variation of intensity only. The subject can

eliminate all sensation of flicker by balancing the intensities of the two lights making them equiluminant. (In fact, luminance is operationally defined using this procedure, known as heterochromatic flicker photometry.) When the intensities are not balanced, the luminance flicker can be seen at frequencies as high as 50-60 Hz. The temporal parameters of the chromatic system are easily incorporated into the three-dimensional halftoning algorithms described in section 2.2.

4. Summary

A number of methods are available by which traditional halftoning algorithms may be tuned to visual system properties. Often this can be done at no additional computational cost merely by careful selection of algorithm parameters. The realizable benefits depend to a large extent on the relation between the display resolution and final viewing conditions.

5. Acknowledgements

This work was supported by NASA RTOP's 505-64-53 and 506-59-65, and DARPA's High Resolution Systems Program.

References

1. Ulichney, R., *Digital Halftoning*, MIT Press, Cambridge, Mass., 1987.
2. Robson, J. G., "Spatial and temporal contrast sensitivity functions of the visual system," *J. Opt. Soc. Am.*, vol. 56, pp. 1141-1142, 1966.
3. van Nes, F. L., Koenderink, J. J., Nas, H., and Bouman, M. A., "Spatiotemporal modulation transfer in the human eye," *J. Opt. Soc. Am.*, vol. 57, pp. 1082-1088, 1967.
4. Mitsa, T. and Parker, K. J., "Digital halftoning technique using a blue noise mask," *J. Opt. Soc. Am. A*, pp. 1920-1929, 1992.
5. Floyd, R. and Steinberg, L., "An adaptive algorithm for spatial gray scale," *SID 1975 Symp. Dig. Tech. Papers*, p. 36, 1975.
6. Carnevali, P., Coletti, L., and Patarnello, S., "Image Processing by Simulated Annealing," *IBM J. Res. Develop.*, vol. 29, pp. 569-579, 1985.
7. Anastassiou, D., "Neural net based digital halftoning of images," *Proc. IEEE Symp. Circuits & Systems*, pp. 507-510, 1988.
8. Analoui, M. and Allebach, J. P., "Model based halftoning using direct binary search," in *Human Vision, Visual Processing and Digital Display III*, ed. B. E. Rogowitz, J. P. Allebach, and S. A. Klein, pp. 96-108, Proc. SPIE 1666, paper #9, 1992.
9. Mulligan, J. B. and Ahumada, A. J. Jr., "Principled halftoning based on human vision models," in *Human Vision, Visual Processing and Digital Display III*, ed. B. E. Rogowitz, J. P. Allebach, and S. A. Klein, pp. 109-121, Proc. SPIE 1666, paper #10, 1992.
10. Balasubramanian, R., Carrara, D., and Allebach, J., "Color image quantization using direct binary search," in *Human Vision, Visual Processing and Digital Display IV*, ed. B. E. Rogowitz and J. P. Allebach, Proc. SPIE 1913, paper #20, 1993.
11. Mulligan, J. B. and Ahumada, A. J. Jr., "Principled methods for color dithering based on models of the human visual system," *SID Int. Symp. Dig. Tech. Papers*, pp. 194-197, 1992.
12. Kelly, D. H., "Spatio-temporal frequency characteristics of color-vision mechanisms," *J. Opt. Soc. Am.*, vol. 64, pp. 983-990, 1974.
13. Koenderink, J. J. and van Doorn, A. J., "Spatiotemporal contrast detection threshold surface is bimodal," *Optics Letters*, vol. 4, pp. 32-34, 1979.
14. Noorlander, C., Heuts, M. J. G., and Koenderink, J. J., "Sensitivity to spatiotemporal combined luminance and chromaticity contrast," *J. Opt. Soc. Am.*, vol. 71, pp. 453-459, 1981.
15. Noorlander, C., Koenderink, J. J., den Ouden, R. J., and Edens, B. W., "Sensitivity to spatiotemporal colour contrast in the peripheral visual field," *Vision Res.*, vol. 23, pp. 1-11, 1983.

16. Watson, A. B., "Temporal sensitivity," in *Handbook of perception and human performance, vol. I Sensory processes and perception*, ed. K. R. Boff, L. Kaufman, and J. P. Thomas, John Wiley and Sons, New York, 1986.
17. Mulligan, J. B., "Methods for spatiotemporal dithering," *SID Int. Symp. Dig. Tech. Papers*, pp. 155-158, 1993.
18. Mulligan, J. B., "Digital halftoning methods for selectively partitioning error into achromatic and chromatic channels," in *Human Vision and Electronic Imaging: Models, Methods, and Applications*, ed. B. E. Rogowitz and J. P. Allebach, pp. 261-270, Proc. SPIE 1249, paper #21, 1990.

